Manuscript version: Published Version

The version presented in WRAP is the published version (Version of Record).

Persistent WRAP URL:

http://wrap.warwick.ac.uk/142724

How to cite:

The repository item page linked to above, will contain details on accessing citation guidance from the publisher.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

warwick.ac.uk/lib-publications

Does the Risk of Poverty Reduce Happiness?

STEFANO A. CARIA
University of Oxford

PAOLO FALCO
Organization for Economic Cooperation and Development
and University of Oxford

I. Introduction

Economic outcomes are often characterized by pervasive uncertainty. This is particularly true when insurance markets and safety nets are incomplete, saving opportunities are limited, and many individuals rely on risky entrepreneurial activities to generate their incomes (Banerjee and Duflo 2007). Poverty, which constitutes a possible outcome for many, has in turn profound impacts on the quality of the lives people live. Recent studies have shown that low income correlates with lower life satisfaction and with a larger loss in well-being after shocks in other domains of life (Clark, Frijters, and Shields 2008; Kahneman and Deaton 2010). In developing countries, where poverty is widespread, the correlation between economic outcomes and life satisfaction is even stronger (Howell and Howell 2008).

In this study we investigate the relationship between the risk of income poverty and life satisfaction (interpreted hereafter as "happiness") and the link between sensitivity to downward risks and decision-making.[1] In particular, we tackle the following two questions. Is there a connection between happiness and the risk

[1] Researchers distinguish two components of happiness (Kahneman and Deaton 2010). The first component is life satisfaction, the evaluation we make of our own life. The second component is emotional well-being, or the tendency to experience positive or negative affect. In this paper we analyze responses from a survey question on life satisfaction, and hence we focus the analysis on the first component. Throughout the text, we use the terms "happiness" and "life satisfaction" interchangeably.

of poverty? And how are people's decisions affected by exposure to such risk?[2] While the connection between life satisfaction and low income has been heavily researched, the one between life satisfaction and the risk of income poverty is still unexplored. This is partly due to the challenges of estimating the probability distribution of income convincingly. At the same time, it appears to be a very important area of research, especially in developing countries, where widespread exposure to uninsured shocks makes the risk of future income poverty pervasive for both poor and nonpoor households.[3] Evidence is also missing on the connection between the determinants of happiness and those of decision-making: are the same individuals whose happiness is sensitive to downside risk loss averse in economic decisions? Vulnerability may affect individual behavior in ways that are detrimental to economic efficiency.[4] Such evidence is thus a necessary first step toward a full assessment of the welfare effects of economic vulnerability.

The context of our analysis is the urban labor market in Ghana, a growing African country. Ghana is an interesting setting for our analysis, as the country experienced substantial poverty reduction in recent years (Nsowah-Nuamah, Teal, and Awoonor-Williams 2010) while, as our results suggest, large numbers are still exposed to a significant risk of poverty. Given the novelty of the question and of the testing strategy, our results provide leads that may prove relevant in other contexts as well.

Our estimate of the risk of poverty builds upon work by Chaudhuri (2003) and Chaudhuri, Jalan, and Suryahadi (2002), who propose two indexes of vulnerability to poverty that are amenable to empirical estimation based on panel and cross-sectional variation, respectively.[5] Using data from the Ghana Urban Household Panel Survey, a long panel data set gathered by the Centre for the Study of African Economies (CSAE) in urban Ghana, we obtain estimates of the

---

[2] We choose to focus on income, rather than consumption, for two main reasons. First, we aim to link directly to the existing literature on subjective well-being, which has widely explored the relationship between happiness and income. Second, in urban contexts like the one we are studying, labor earnings are typically the main source of income (for this reason, throughout the article the words "income" and "earnings" are used interchangeably), and earning shocks are directly transmitted to consumption. Changes in income and consumption are hence likely to be tightly correlated. We fully acknowledge the difficulties that arise from attempting to measure income precisely in a low-income setting with widespread informality, such as Ghana, especially when self-employment is prominent. Sec. II provides a detailed discussion of the methodological challenges involved.

[3] For example, in a recent study of seven West African capitals, Bocquier, Nordman, and Vescovo (2010) construct a multidimensional index of employment vulnerability and find that 85% of private-sector workers are vulnerable on the basis of at least one criterion in 2002–3.

[4] Throughout the analysis the term "vulnerability" is used to refer to the risk of falling below the income poverty line.

[5] The two indexes are reviewed in a survey article by Ligon and Schechter (2004), who compare the performance of different vulnerability measures through Monte Carlo simulations.

two indexes for a representative sample of working-age Ghanaian earners. We focus more extensively on the panel measure, since it enables us to estimate individual-specific vulnerability. We further rely on the longitudinal nature of the data to investigate the relationships of interest, between the risk of income poverty and life satisfaction. Improving upon most of the existing literature on happiness in developing countries, we are able to control for individual fixed effects in the happiness model, ruling out potential biases from unobserved personality traits. Previous studies have indeed highlighted the importance of unobserved heterogeneity in happiness regressions (Ferrer-i-Carbonell and Frijters 2004; Graham, Eggers, and Sukhtankar 2004; Powdthavee 2010).

Our main result is a strong negative relationship between vulnerability to income poverty and workers' happiness, over and above the positive income effect commonly documented in the literature. It is both statistically significant and economically meaningful. Reducing the risk of poverty by 20 percentage points (which amounts to entirely offsetting the risk of poverty for the median worker in our sample) has the same effect on well-being as increasing earnings by 50%. When we bootstrap the estimation sequence to account for imprecision in the measure of vulnerability, the results do not change. Upon testing for the role of two-sided uncertainty, as opposed to downward income losses, we find that the effect of downward vulnerability on happiness is more evident. These findings become more compelling when we consider the extent of the vulnerability to poverty we uncover. About 35% of all workers, and 15% of currently nonpoor workers, face a risk of poverty of at least 50%. Vulnerability decreases the life satisfaction of a large pool of individuals.

In addition, we analyze the choices of a subsample of respondents in a set of behavioral games designed to elicit attitudes toward risky prospects. Our maximum likelihood estimates reveal that subjects are characterized, on average, by a substantial degree of loss aversion. We are careful not to collapse the distinct notions of experienced utility and decision utility (Kahneman, Wakker, and Sarin 1997).[6] Our findings from the behavioral experiment show that, besides influencing subjective well-being, downside risk also has an appreciable impact on economic decisions.

Our work relates to two different strands of the literature. First, we contribute to the study of downside risk in developing countries. This literature has focused on measurement (Chaudhuri 2003; Ligon and Schechter 2004), the persistence of downside shocks (Dercon 2004; Dercon, Hoddinott, and Wol-

---

[6] "Decision utility" refers to the weights that people assign to outcomes when making choices. "Experienced utility" refers to the quality of an experience. The life-satisfaction question we employ in this analysis asks respondents to make an evaluation of the latter. The behavioral experiments allow us to make inferences about the former.

dehanna 2005), and the strategies employed to minimize and cope with shocks (Rosenzweig and Binswanger 1993; Dercon 1996; Fafchamps 2003, 2009; Dercon and Christiaensen 2007). Most importantly, our analysis contributes to the growing literature on the determinants of happiness. A number of empirical papers have documented a cross-sectional correlation between income and happiness (Kahneman and Deaton 2010), which does not disappear once individual fixed effects are accounted for (Ferrer-i-Carbonell and Frijters 2004; Powdthavee 2010). A separate concern has been that of adaptation: the happiness effects of income gains seem transitory and tend to disappear once income reference points have adjusted (Easterlin 2001; Frey and Stutzer 2002; Di Tella, Haisken-De New, and MacCulloch 2007; Knight and Gunatilaka 2008). The literature has also explored the effect of social comparisons on well-being (Blanchflower and Oswald 2002; Kingdon and Knight 2004; Luttmer 2005). Our contribution is to highlight the fact that risk—in particular, the risk of poverty—is a major negative determinant of life satisfaction. Moreover, we show that the same people who manifest loss sensitivity in life evaluation make economic decisions consistent with loss aversion.

The results of this analysis bear important policy implications that may generalize well beyond the African context. Our findings provide clear motivation for policy interventions to reduce people's exposure to (downside) risk. They also suggest that non-Rawlsian models of growth, whereby "someone may be left behind," may fail to enhance general welfare despite rising average incomes, if the risk of falling behind is sufficiently widespread. Finally, loss aversion motivates individuals to forgo economic opportunities that are profitable in expectation but may involve outcomes below the reference point. A reduction in vulnerability may result in efficiency gains too.

The article is structured as follows. Section II introduces the data we use in the analysis. Section III outlines the empirical strategy. First, it explains the methodology to estimate income vulnerability; second, it outlines the happiness model. Section IV presents and discusses the results. Section V concludes.

II. Data
Our analysis is based on data from the Ghana Urban Household Panel Survey (GUHPS), which was conducted by the CSAE in the cities of Accra, Kumasi, Takoradi, and Cape Coast starting in 2004.[7] Respondents were drawn by stratified random sampling of urban households from the Population and Housing Census of 2000. The survey was designed to cover all household members of

[7] The first wave was collected between the end of 2003 and the beginning of 2004, but for simplicity we refer to it as 2004.

working age at the time of the interview. This paper focuses on the period between 2004 and 2009, when the survey was repeated every year with the sole exception of 2007.[8] Panel data sets of this length are unusual in developing countries and are particularly uncommon in Africa.[9]

A module on subjective well-being—designed in accordance with the existing literature—was added to the survey in 2005. Our analysis focuses on the answers to the following two questions: (1) "All things considered, how satisfied are you with your life as a whole these days?" and (2) "All things considered, how satisfied are you with your current work?" In both cases, the options given to respondents were "1. Very Dissatisfied, 2. Dissatisfied, 3. Neither Satisfied Nor Dissatisfied, 4. Satisfied, 5. Very Satisfied." Figure 1 depicts the distribution of answers. Responses appear to be skewed toward positive values. For our quantitative analysis, we attribute numerical values on a scale from 1 to 5 to these answers, where 1 corresponds to "Very Dissatisfied" and 5 to "Very Satisfied." Despite early criticism of their ability to accurately capture well-being (e.g., Mullainathan and Bertrand 2001), these measures have been consistently used throughout the literature. Moreover, psychologists have recently been able to validate the use of these questions, by showing their correlation with other measures of well-being, such as smiling more frequently (Graham, Eggers, and Sukhtankar 2004; Layard 2005; Oswald and Wu 2010).

A selection of key summary statistics for our sample of interest is presented in table 1. The average worker is 35.3 years old and has 8 years of formal education. Most of them are self-employed (largely in the informal sector), as is typical in the labor markets of many developing countries.[10] For the self-employed,

---

[8] Subsequent survey waves were unavailable to us at the time of writing. Moreover, extending the analysis beyond 2009 would pose severe challenges due to a change in the survey question on life satisfaction introduced in 2010.

[9] The panel is unbalanced, but attrition is not an absorbing state, in the sense that respondents who are not interviewed in a given wave are kept in the sample and reinterviewed in subsequent years. Out of the initial sample of respondents interviewed in 2004, over 92% were reinterviewed at least once in the following years, and about 65% were observed in at least three waves. Random attrition would decrease precision and pose a classical problem of attenuation bias in the happiness model. In this case, the large effects we estimate would be a lower bound of the true impacts. On the other hand, if the people whose happiness is least affected by the risk of poverty are most likely to drop out, our estimates may be biased upward. As a robustness check, the analysis presented below has been repeated on the strictly balanced sample of individuals who are interviewed in all waves, and we find that this does not affect the results. The point estimates of the coefficients discussed below do not change substantially, despite a drop in precision due to the fall in sample size. Previous studies on these data have found no evidence of selective attrition biasing the results (see Falco et al. 2011, 2015).

[10] The sample described in table 1 is confined to workers for whom self-reported subjective well-being data are available. Furthermore, it includes only paid workers for whom income is observed and we are therefore able to construct a measure of income vulnerability.
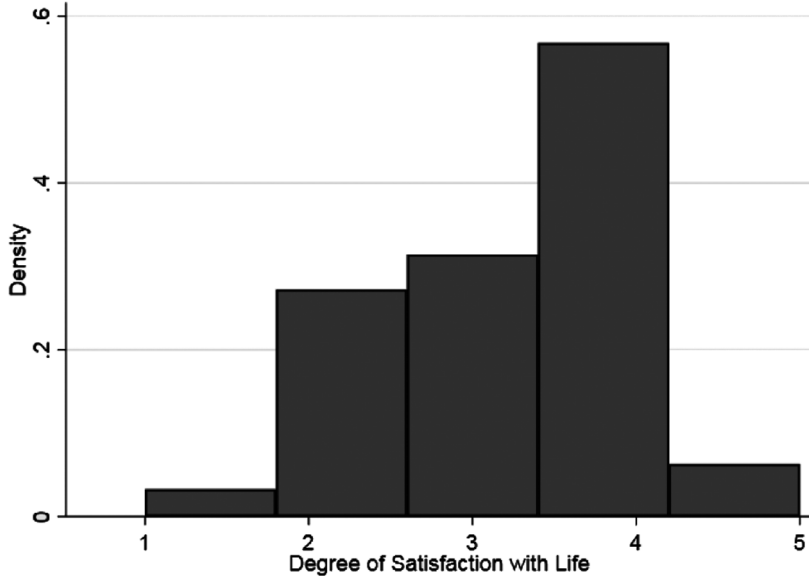
Figure 1. Distribution of life satisfaction across the population (*top*) and by level of income (employed people only; *bottom*).

TABLE 1
SUMMARY STATISTICS, 2005

Variable	Mean	Standard Deviation
Age	35.27	10.73
Education (years)	8.038	4.032
Male	.446	.497
Private-sector wage	.265	.442
Public-sector wage	.074	.262
Ln(employees)	.133	.418
Ln(firmsize)	.842	1.551
Years since started current job	10.625	9.888
Married	.53	.499
Ethnicity:		
Ga-Dangme	.145	.352
Ewe	.06	.237
Mole-Dagbani and Hausa	.096	.295
Other ethnicity	.112	.316

Note. Restricted to observations in the happiness model, pooling survey waves. There were 740 observations, of which 738 include ethnic origin.

income is calculated as revenue minus cost (i.e., profit).[11] In low-income settings, such as urban Ghana, it is difficult to obtain precise estimates of such profit, and we are aware of the challenge. A number of published studies, however, have successfully relied on profit information from GUHPS (see Falco and Haywood 2016 and Falco et al. 2015 for the most recent examples). Moreover, a study by Fafchamps et al. (2012) lends additional support to the data. They use personal digital assistant (PDA) technology to cross-check profit figures in a sample of microentrepreneurs in urban Ghana that is very similar to the GUHPS sample. They find that only a small proportion of profit calculations are imprecise.[12]

---

[11] The resulting profits are attributed to the owner of the business. The remaining household members who work in the business are categorized as paid or unpaid family labor, depending on whether they receive an income. This may lead to some imprecision, as it may be difficult to clearly identify business ownership in some cases. However, the fact that the large majority of the self-employed in our sample (about 80%) does not employ any other person (including household members, either paid or unpaid) makes the potential magnitude of the problem relatively small. When a self-employed person works by himself/herself (i.e., hires no other workers), the value of the variable "employees" is set to 1. Consequently, its log in the regressions below equals 0. The same approach was followed by Falco et al. (2011).

[12] There might be additional reasons of concern beyond the precision of profit calculations. For instance, in the absence of accurate bookkeeping, microentrepreneurs may simply recall the "wrong" revenue and cost figures, which might equally affect standard (paper-based) and PDA-based surveys if the latter are unable to spot inconsistencies across revenue, costs, and profit figures. Most microenterprise surveys suffer from such problems, and the fact that GUHPS data have already been successfully employed in several published studies lends credibility to our analysis.

III. Empirical Methodology

A. Constructing an Indicator of Vulnerability to Income Poverty

This section outlines the methodology to construct the vulnerability indicators used in the remainder of the analysis. For a detailed discussion of the relative merits of different vulnerability indexes, the reader is referred to the survey paper by Ligon and Schechter (2004). The analysis in this article draws on the two measures proposed by Chaudhuri (2003) and Chaudhuri, Jalan, and Suryahadi (2002). The former relies on time-series variation in individual earnings and suits particularly well the characteristics of our data set, where subjective well-being is recorded for the same individuals over a number of consecutive years (in addition to income and other worker characteristics from which we can model vulnerability). The latter method attempts to model cross-sectional variation in earnings and to infer from it the degree of individual vulnerability. Ligon and Schechter (2004) compare the performance of these two (and several other) vulnerability indexes via Monte Carlo simulations, and their conclusion is in favor of the panel approach as the best-performing indicator of actual vulnerability.[13] In light of their findings and of the data at our disposal, our main focus is on the panel method. The appendix (available online) presents alternative estimates based on the cross-sectional approach and offers insights on the differences between the two.

   Following Chaudhuri (2003), the income vulnerability of a worker at time t is defined as the probability that the worker's income will fall below a certain threshold z in the next period. This differs from a simple measure of volatility. The distribution from which a given worker's income is drawn can have very large volatility, but if most of the variation occurs above the threshold z, this worker has a low probability of earning less than z and hence has low vulnerability.

   Let $v_{i,t}$ be the inverse of vulnerability, that is, i's probability at t of earning an income above z at $t + 1$:

$$v_{i,t} = \Pr(y_{i,t+1} > z). \tag{1}$$

When z is the poverty line, $v_{i,t}$ is the probability at t that household i will not be income poor at $t + 1$. Following standard Mincerian earnings analysis, we assume that income is generated by the following process:

---

[13] This is a sensible conclusion, considering the likely presence of unobserved individual fixed effects that cannot be controlled for in a cross-sectional model of earnings and could therefore mislead the analysis of vulnerability. An interesting alley for future research would be to explore how the results of this analysis would change upon use of other vulnerability measures, including Ligon and Shechter's own index of vulnerability (see Ligon and Schechter 2003).

$$\ln(y_{i,t}) = \delta X_{i,t} + \eta_i + \tau_t + e_{i,t}, \tag{2}$$

where $X_{i,t}$ is a bundle of observable characteristics, η_i is an individual unobservable fixed effect, τ_t captures time effects that are common across workers (e.g., aggregate income growth factors and common shocks), and $e_{i,t}$ is a stochastic component.

Second, we assume the variance of $e_{i,t}$ to be a function of worker and household characteristics,

$$\ln\left(\sigma^2_{\ln y_{i,t}}\right) = \theta K_{i,t} + \xi_i, \tag{3}$$

where ξ_i is an individual fixed effect in the model of income variance and $K_{i,t}$ may or may not contain additional worker characteristics outside the set $X_{i,t}$. In Chaudhuri (2003), $K_{i,t}$ and $X_{i,t}$ coincide.

The variance of the stochastic component can be modeled empirically with the log of first-stage residuals from the earnings model:

$$\ln(\hat{e}^2_{i,t}) = \theta K_{i,t} + \xi_i + \omega_{i,t}, \tag{4}$$

given that

$$\frac{1}{T}\sum_{t=1}^{T}\hat{e}^2_{i,t}\rightarrow_p \sigma^2_{\ln y_{i,t}}. \tag{5}$$

Assuming income to be (log)normally distributed and Φ to be the cumulative distribution function of the lognormal distribution, we can now compute the probability of not earning less than z at $t + 1$ for every worker i by using the following expression:

$$\hat{v}_{i,t} = \widehat{\Pr}\left(\ln(y_{i,t}) > \ln(z)|X_{i,t}, K_{i,t}, \hat{\delta}, \hat{\theta}\right) = 1 - \Phi\left(\frac{\ln(z) - \hat{\mu}_{i,t}}{\hat{\sigma}_{i,t}}\right), \tag{6}$$

where $\hat{\mu}_{i,t}$ denotes the predicted value of (log) income and $\hat{\sigma}_{i,t}$ the predicted variance.[14]

This measure does not differentiate transitory from permanent shocks. When saving, credit, and insurance options are limited, both transitory and permanent shocks can have a significant impact on consumption. However, in the presence of an effective saving technology, the two may have different impacts on welfare,

---

[14] The reader should note that, unlike the definition in eq. (1), our estimates of vulnerability are obtained as the probability of falling below the poverty line, given worker characteristics at t, rather than $t + 1$. This choice was based on the idea that workers are most likely to assess their future prospects on the basis of their current characteristics, some of which might themselves be stochastic and subject to unpredictability.

since transitory fluctuations can be smoothed out through precautionary savings. We acknowledge that separating transitory and permanent shocks would be useful, especially in the largest urban areas, where financial markets are likely to be more accessible. Unfortunately, attempting to explicitly separate the permanent from the transitory component of income variation (e.g., following the approach by Meghir and Pistaferri 2004) would pose major challenges, given the length of the panel at our disposal. This remains an open alley for future research.

B. Empirical Model of Happiness

Having constructed a measure of income vulnerability, we can now explore its relationship with subjective well-being. The following equation describes our workhorse model of happiness:

$$h_{i,t} = \beta y_{i,t} + \gamma v_{i,t} + \delta Z_{i,t} + \kappa_i + \varepsilon_{i,t}, \tag{7}$$

where $h_{i,t}$ is worker i's level of life satisfaction in period t, $y_{i,t}$ is income at time t, and $v_{i,t}$ is the inverse of vulnerability in the same period, $Z_{i,t}$ is a vector of worker characteristics that are expected to be correlated with life satisfaction, and κ_i is an unobserved happiness fixed effect that accounts for unobserved traits that make an individual naturally more (or less) prone to be satisfied with his/her life (e.g., optimism). Our main hypothesis is that β and γ are positive (once again, note that $v_{i,t}$ is the inverse of vulnerability and hence a "good" in this specification): increasing income and decreasing vulnerability enhance life satisfaction. In order to test it, we attempt to overcome several identification challenges.

First, a number of time-varying and time-invariant determinants of happiness may be correlated with income and vulnerability. If omitted from the analysis, those variables may bias the results. Among the time-invariant factors, one can think of personality traits and endowments of social and human capital (which may have direct impacts on both job prospects and life satisfaction). More extroverted and optimistic individuals, for instance, may be both "naturally" satisfied with their life and more likely to find good, secure employment or, equally plausibly, more willing to face the risks and uncertainty of entrepreneurship. The same may hold for educated or well-connected people. Among the time-varying unobservables, working conditions are a first, obvious source of bias. Powdthavee (2010) argues that income gains are often correlated with deterioration in the conditions of work, and the latter may have an important influence on life satisfaction. Vulnerability might also be correlated with working conditions, although we have no strong a priori evidence of the sign of such correlation. Relative income is a third potentially confounding factor. Exten-

sive empirical evidence has been generated showing that relative income is correlated with the life satisfaction of individuals in both developed and developing countries (Blanchflower and Oswald 2002; Kingdon and Knight 2004; Luttmer 2005), and it is natural to assume that relative income will be correlated with absolute income and vulnerability. We attempt to account for these potential sources of bias by including in the model controls for working conditions (proxied by a measure of satisfaction with work) and for a worker's position in the income distribution. Most importantly, thanks to our panel data set we are able to control for all time-invariant individual characteristics correlated with happiness (e.g., personality traits).

The second challenge is methodological: life satisfaction is generally recorded in data sets like GUHPS as a categorical variable. Modeling it as a discrete (ordered) outcome would, therefore, appear to be the most appropriate approach. However, such an approach would not easily lend itself to controlling for those time-invariant unobservables that we have argued are of great relevance in the determination of life satisfaction. To address this issue, Ferrer-i-Carbonell and Frijters (2004) develop a conditional estimator for the fixed-effects logit model. Their findings show that "assuming ordinality or cardinality of happiness scores makes little difference, while allowing for fixed-effects does change results substantially" (Ferrer-i-Carbonell and Frijters 2004, 641).[15] It therefore seems justifiable to assume cardinality of the life-satisfaction indicator and use the corresponding estimators.

Third, issues of reverse causality may arise in the analysis. High levels of life satisfaction may help individuals earn higher incomes or reduce their income vulnerability (Graham, Eggers, and Sukhtankar 2004; De Neve and Oswald 2012). Such effects may again bias the estimated coefficients β and γ. In order to fully address this problem, we would be required to specify an FE-IV (fixed effects–instrumental variable) regression approach. However, doubts are often raised about the validity of the instruments proposed by authors who have attempted the IV or FE-IV approach for income, such as Knight and Gunatilaka (2008) and Powdthavee (2010).[16] Hence, we do not attempt to instrument vulnerability, while fully acknowledging the possibility that these concerns might be important.

[15] Ferrer-i-Carbonell and Frijters (2004) base their analysis on data from the German Socio-economic Panel. The life-satisfaction question in their data uses the same phrasing as that in GUHPS. The only difference is that it is posed on a 1–10 scale, while in Ghana the scale is from 1 to 5. This should not pose a challenge for comparability. The results reported in our paper confirm the main conclusion of Ferrer-i-Carbonell and Frijters (2004) on the importance of controlling for individual fixed effects in empirical models of life satisfaction.
[16] Furthermore, the vulnerability variable has been constructed as a deterministic function of the predicted values of an earnings model, which would complicate an IV strategy.

Finally, the vulnerability index is a nonlinear function of the first two moments of the earnings distribution, which are both modeled as functions of household and individual characteristics in the first stage of the estimation. It follows that the happiness model (where we include both income and vulnerability on the right-hand side) contains two functions of those characteristics among the regressors. Separate identification of these two functions implicitly relies on assumptions regarding the relationship between income and well-being. Existing studies have often imposed linearity on the relationship, and, for comparability, we choose the same approach.[17]

IV. Results

We present here three sets of results. First, we discuss our estimates of income vulnerability. Second, we present a number of regressions of happiness on vulnerability, which constitute the central results of our analysis. This section also offers a test to distinguish between the effect of vulnerability and that of two-sided uncertainty. Third, we show the results of a complementary analysis of workers' attitudes to gains and losses based on data from a field experiment.

A. Vulnerability Estimates

Table 2 shows the results from estimating the earnings and variance models used to predict vulnerability later in the analysis. The first feature of the results is that while the earnings model (col. 1) shows a relatively high predictive power, trying to predict the variance of earnings proves to be a much more challenging exercise (cols. 2–5).[18]

Upon experimenting with different earnings specifications, we concluded that the best model is one that controls for individual fixed effects and for a set of key time-varying covariates (col. 1), the choice of which is grounded in a long-established literature on Mincerian earnings regressions (see Rankin, Sandefur, and Teal 2010 for an application to Ghana using the GUHPS data set). The results confirm a number of standard patterns observed in related studies of earnings in sub-Saharan Africa. First, we find a statistically significant effect of firm size on earnings (captured by positive coefficients on the log of firm size for wage employees and on the log of the number of hired employees for the self-employed). Second, we detect a sizable civil-service premium and a positive premium for longer tenure in the job. Third, while the linear effect

---

[17] Fafchamps and Shilpi (2008, 2009) report nonparametric results that show a linear relationship between consumption expenditures and subjective satisfaction with consumption levels, lending empirical support to this modeling choice.

[18] This is to be expected, partly because a fraction of what appears to be true variation in earnings may in fact be random measurement error.

TABLE 2

ESTIMATION OF VULNERABILITY

Dependent Variable	y (1)	$\sigma^2(K)$ (2)	$\sigma^2(X)$ (3)	$\sigma^2(X^2)$ (4)	$\sigma^2(X, \text{FE})$ (5)
Age		.035	.019		−.0005
		(.028)	(.026)		(.001)
Age^2	−.0003	−.0004	−.0002	−5.49E−08	−.0005
	(.0004)	(.0003)	(.0003)	(4.76E−08)	(.001)
Education		−.028	−.055*		
		(.034)	(.032)		
Education^2		.003	.004*	7.17E−06	
		(.002)	(.002)	(8.15E−06)	
Male		.134	.133	.127	
		(.095)	(.092)	(.092)	
Private-sector wage	−.147**	−.991***	−.944***	−.967***	−.234
	(.074)	(.149)	(.145)	(.121)	(.177)
Public-sector wage	.196*	−1.342***	−1.319***	−1.313***	−.120
	(.112)	(.184)	(.180)	(.177)	(.269)
Ln(employees)	.187***	.008	.040	.044	−.030
	(.051)	(.115)	(.113)	(.057)	(.121)
Ln(firmsize)	.055**	.005	−.013	−.0003	.062
	(.021)	(.042)	(.042)	(.008)	(.051)
Years since current-job start	.005*	.003	.004	.0002	.005
	(.003)	(.006)	(.006)	(.0002)	(.007)
Married		−.052			
		(.095)			
Ethnicity:					
Ga-Dangme		−.037			
		(.119)			
Ewe		.392**			
		(.171)			
Mole-Dagbani and Hausa		.578***			
		(.155)			
Other		−.181			
		(.162)			
Constant	2.588*	−3.041***	−2.642***	−2.514***	−3.293**
	(1.346)	(.519)	(.491)	(.190)	(1.294)
Individual fixed effects	Yes				Yes
Time dummy	Yes	Yes	Yes	Yes	Yes
Observations[a]	3,659	3,014	3,110	3,110	3,110
R^2	.685	.073	.065	.064	.627

Note. Robust standard errors are in parentheses. y = log(real monthly earnings); $\sigma^2(K)$ = log(variance of y) modeled as a function of K; $\sigma^2(X)$ = log(variance of y) modeled as a function of X; $\sigma^2(X^2)$ = log(variance of y) modeled as a function of X^2; $\sigma^2(X, \text{FE})$ = log(variance of y) modeled as a function of X including individual fixed effects (this specification is used to compute vulnerability subsequently in the paper). X is the set of key regressors in the income model, K is an augmented set of regressors to include potential determinants of the variance. The omitted occupational category is "self-employed"; the omitted ethnicity is Akan. "Public-sector wage" includes all salaried workers in the public sector, including civil servants and workers in public enterprises.

[a] $N \times T$; columns 2–5 have fewer observations than column 1 because the variance analysis must be confined to respondents who appear multiple times in the panel. Column 2 has fewer observations than columns 3–5 because of missing values in the marriage and ethnicity variables.

* Confidence: 90%.

** Confidence: 95%.

*** Confidence: 99%.

of age cannot be estimated when time trends are also controlled for, we find some indication of the typical concavity of the age-earnings profile (though the coefficient on age squared is not statistically significant). Since the estimation in column 1 is carried out with controls for individual fixed effects, it is not possible to separately identify the coefficients on time-invariant characteristics such as education and gender.

We estimate several variance models. We first estimate a model with the same covariates used in the earnings regression and no fixed effects. This model is able to explain only a small portion of the variation in the data. We then estimate a model that includes additional covariates that were not part of the earnings model. Motivated by the observation that social networks can provide an important buffer against negative income shocks, such additional variables include the respondent's ethnic background and marital status. This strategy allows us to marginally increase the fit of the model but requires us to make the assumption that ethnicity and marital status determine only the variance of income and not its mean. Our preferred model is thus the model reported in column 5, where we include the same covariates used in the earnings regression as well as allow for individual fixed effects to determine the variance of income.[19] This strategy is close to the one employed in Chaudhuri (2003), and the measure of vulnerability used in the rest of the paper is based on this method.

Given our estimates of the earnings model, before we can calculate vulnerability we need to define a low-earnings threshold (alternatively referred to as the "poverty line"), z. Figure 2 shows the percentage of people who are below different income thresholds, while figure 3 shows the resulting cumulative distribution of the vulnerability index. Our choice for the remainder of the paper is to set $z = 10$ Ghana cedis per month, which approximately translates into US$40.[20] When we experimented with alternative lines in the vicinity of this value, the main patterns in our results did not change.[21] For $z = $ US$40, the

[19] Ethnicity does not vary over time, and we also have very little variation in marital status. This gives us a second reason to omit ethnicity and marital status in the variance model with fixed effects.

[20] All income figures have been deflated, and the entire analysis is conducted at constant prices with 1997 as the base year. The reader should be alert to the fact that in 2007 the Ghana cedi was converted into the new Ghana cedi at a rate of 10,000 Ghana cedis to 1 new Ghana cedi. For simplicity, all the income figures used in the analysis have been converted into new Ghana cedis. The exchange rate with the US dollar is not adjusted for purchasing-power parity (PPP).

[21] The poverty line we have chosen (10 Ghana cedis per month) translates into nearly US$170, adjusted for PPP. Assuming that the typical earner in our sample sustains a family of five people with one additional adult (older than 14) and three children (or alternatively, a family of four with two additional adults and one child, based on a standard equivalence scale), this translates into about $1.70 per capita a day, which is close to standard definitions of poverty. This figure is also consistent with the responses to a question added to GUHPS in 2009, which asked respondents to report the income level they deemed sufficient to cater to (1) basic needs and (2) a comfortable life. We interpret
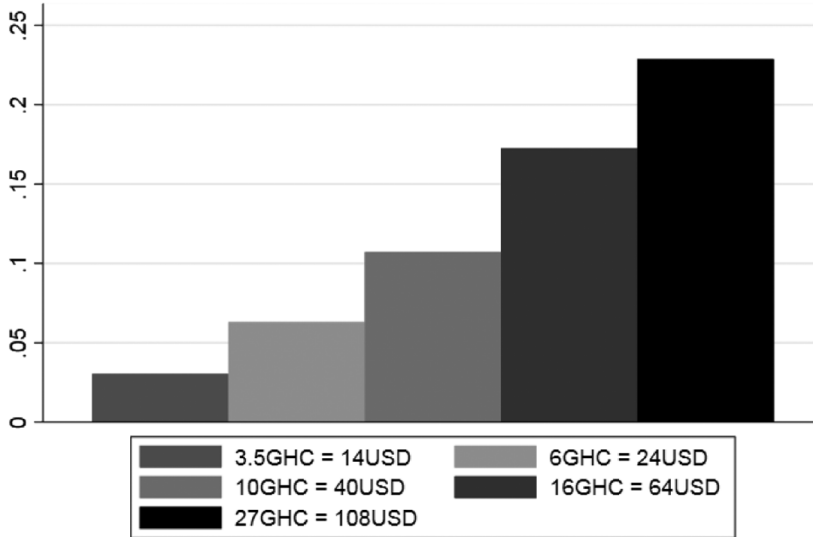
Figure 2. Percentage of employed with income below the low-earnings threshold ($y < z$). GHC = Ghana cedis; USD = US dollars.

risk of poverty we estimate is substantial for large portions of our sample. The central line in figure 3 shows that, on average, 35% of workers in a given year have a chance of 50% or more of being poor in the next period. When we disaggregate the estimation between currently poor and nonpoor workers (fig. 4), we find that even among the latter a sizable group faces a large probability of becoming poor.

B. Happiness

This section presents the results from estimating our happiness model. Figure 1 (*bottom*) plots the histogram of happiness responses after the sample was split by low/high income relative to the poverty line. The histogram shows prima facie evidence of the link between income and happiness that we are attempting

---

the answers as a direct, albeit crude, measure of workers' reference points, and variation below the reference point can be considered as downside risk. Upon plotting those answers (available upon request), we found that for the vast majority of the sample (more than 90%), our chosen poverty threshold lies below both measures of minimum desirable income. This lends strong support to the assertion that our low-income range is within the domain of poverty as perceived by urban Ghanaians. Nonetheless, we are aware that, while we have good reasons to set the poverty line at 10 Ghana cedis per month, other thresholds could be used. We thus experimented with poverty lines that range from about one-third to about three times our preferred value and found no major changes in our main results.

Figure 3. Cumulative distribution of vulnerability for different poverty lines (z_i). GHC = Ghana cedis; USD = US dollars.


Figure 4. Cumulative distribution of vulnerability by current poverty status; $z_t = 10$ (1997) Ghana cedis.

to formally test, with people who are above the low-income threshold more likely to report being "satisfied" with their life.

Table 3 reports the results from estimating the workhorse model of happiness (eq. [7]), first using ordinary least squares (OLS; cols. 1–2) and then controlling for fixed effects (cols. 3–4). Our first result is a positive and significant effect of absolute income on life satisfaction, in line with the existing literature (e.g., De Neve and Cooper 1998). This relationship is evident in the OLS regressions, and, rather strikingly, it does not change significantly once we control for fixed effects. It appears, therefore, that time-invariant unobservables correlated with earnings are not biasing the estimated effect of income on happiness. Interestingly, the size of the estimated coefficient on the log of income in column 3, 0.017, is remarkably close to that estimated by Powdthavee (2010), using data from the British Households Panel Survey and a fixed-effect estimator, 0.019.

TABLE 3
HAPPINESS AND VULNERABILITY

	Pooled Ordinary Least Squares		Fixed Effects	
	(1)	(2)	(3)	(4)
(1 − Vul)	−.006	.005	.087**	.117***
	(.017)	(.018)	(.036)	(.044)
LnRealEarn	.013**	.033***	.017**	.051***
	(.006)	(.012)	(.008)	(.017)
LnWorkSatis	.618***	.615***	.588***	.587***
	(.014)	(.014)	(.025)	(.025)
Married		.027**		.023
		(.010)		(.020)
Age		−.006*		−.019
		(.003)		(.014)
Age^2		.00007*		.0002
		(.00004)		(.0002)
EarnQuart = 2		−.0009		−.026
		(.018)		(.025)
EarnQuart = 3		−.048**		−.069**
		(.023)		(.034)
EarnQuart = 4		−.060*		−.099**
		(.032)		(.048)
Constant	.426***	.489***	.388***	.784***
	(.019)	(.059)	(.036)	(.273)
Observations[a]	2,507	2,507	2,507	2,507
R^2	.45	.454	.422	.425

Note. Robust standard errors are in parentheses. EarnQuart = earnings quartile.
[a] $N \times T$. The number of observations is lower than that in the earnings model in table 2 because of missing values in the life-satisfaction variable (most notably due to the fact that there was no question on life satisfaction in the very first wave of GUHPS [2004]). We keep 2004 data in the earnings regressions to ensure we model the earnings process as precisely as possible.
* Confidence: 90%.
** Confidence: 95%.
*** Confidence: 99%.

On the other hand, table 3 shows that individual fixed effects play an important role in the relationship between vulnerability and life satisfaction. Once we control for them, we find a strong negative relationship between vulnerability and happiness, over and above the income effect just described (recall that in the regression tables this is reported as a positive relationship between the inverse of vulnerability and happiness). This is the key result in the paper. It is both statistically significant and economically meaningful. Reducing the risk of poverty by 20 percentage points (which amounts to entirely offsetting the risk of poverty for the median worker in our sample) has the same effect on well-being as increasing earnings by 50%.

The change in the coefficient on the inverse of vulnerability, which we obtain upon controlling for fixed effects, is an indication that individuals who tend to be more satisfied with their lives (high κ_i) face more downside risk.[22] This could be the case if a common personality trait determines life satisfaction as well as occupational choice. For example, individuals who value autonomy may be happier as well as more likely to work in sectors such as self-employment, where autonomy comes at the cost of a higher risk of low earnings (Benz and Frey 2008; Falco et al. 2015). We investigate this hypothesis by computing the correlation between the estimated fixed effects from the happiness model in equation (7) and the probability of being self-employed (which in Ghana, as shown by Falco 2014, is associated with higher earnings risk than salaried work). Consistent with our hypothesis, we find a positive and highly significant correlation (not shown for conciseness).

Our estimation also includes controls for work satisfaction (proxying changes in working conditions), earnings quartile, age and its square, and marital status. Work satisfaction is closely correlated with life satisfaction and shows by far the biggest positive coefficient in the life-satisfaction regression.[23] The earnings quartile dummies allow us to control for the position of respondents in the income distribution, which has been shown to be a significant predictor of well-being (Clark, Frijters, and Shields 2008). Their inclusion in the regression model does not affect our main results (if anything, the coefficient on vulnerability increases slightly). This should reassure the reader that changes in relative income are not confounding our estimates of the relationship between income vulnerability and happiness.[24]

---

[22] The fixed effects appear to correlate negatively with the "inverse" of vulnerability. Hence, they correlate positively with downside risk.

[23] As a robustness check, we tried to exclude work satisfaction from the estimation, and the results did not change significantly (we detected only a slight increase in the effect of vulnerability).

[24] In addition to what is reported in the table, we experimented with finer quantile disaggregation (quintiles and deciles), and the main results did not change.

Finally, the vulnerability index has been constructed with estimates from a first-stage model of earnings. Hence, it carries a degree of statistical imprecision that could pose a challenge to the significance of our estimates in the second-stage model of happiness. In order to check the robustness of our results to such a concern, we bootstrap the entire estimation sequence (including the first stage to construct the vulnerability index), sampling with replacement to obtain 200 replications of the original sample. The results are summarized in figures 5 and 6, where we plot the distribution of the bootstrapped coefficients on earnings and on the inverse of vulnerability (from the happiness model), and they are consistent with the discussion so far.

C. *Alternative Explanations*

The measure of vulnerability employed in this paper focuses on the notion of exposure to downside risk, and we interpret our findings as showing that the risk of income poverty has a significant impact on well-being. An alternative explanation could be that individuals dislike income volatility per se (rather than exposure to downside risk). In this section we attempt to disentangle the two hypotheses by replacing the vulnerability measure used so far with two-sided measures of earnings volatility.

First, we use the raw squared residual $\hat{e}_{i,t}^2$ from a first-stage earnings regression with fixed effects as a proxy for income volatility and find no significant
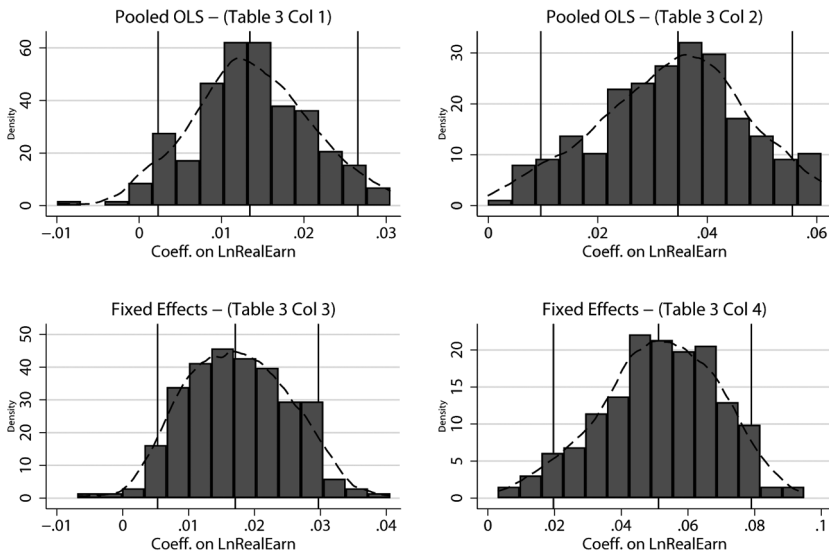

Figure 5. Bootstrapped distribution (200 samples) of the coefficient on LnRealEarn. The central vertical line on each graph indicates the median, and the two outer vertical lines indicate the 5th and 95th percentiles of the distribution. The dashed line represents a kernel density of the distribution.

Figure 6. Bootstrapped distribution (200 samples) of the coefficient on (1 − Vul). The central vertical line on each graph indicates the median, and the two outer vertical lines indicate the 5th and 95th percentiles of the distribution. The dashed line represents a kernel density of the distribution.
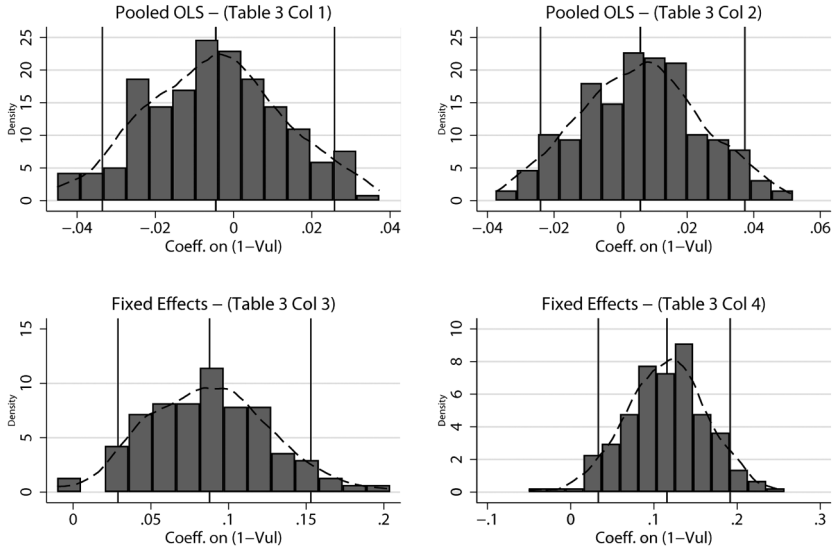
relationship with happiness (table 4), despite the sign of the estimated effect always being negative (as we would expect if workers are risk averse).[25] The lack of statistical significance might be due to the fact that ex post realizations of the shock are a noisy proxy of the expected degree of vulnerability workers perceive (and are affected by). A way to circumvent the problem is to model the variance of these residuals, as we did in Section III.A, and use the predicted value as a measure of expected variance. Upon doing that, we document once again a negative effect of volatility on life satisfaction that is not statistically significant (the results are not reported for conciseness but are available upon request). Overall, this evidence points to the conclusion that vulnerability to downside income risk, as analyzed in the previous section, plays a more clear-cut role in the determination of well-being than two-sided volatility.

D. Choice among Risky Prospects

Our final piece of evidence comes from a behavioral experiment that studies individual choices between risky prospects when downside risk is present and when it is absent. Our objective is to investigate whether downside risk affects the economic decisions of the respondents in our sample. We do so by estimating the level of loss aversion implied by the respondents' choices in a series

[25] Using experimental data, Falco (2014) shows that the majority of GUHPS respondents are indeed risk averse.

TABLE 4
VULNERABLE TO DOWNSIDE RISK OR AVERSE TO UNCERTAINTY? (RESIDUAL)

	OLS1 (1)	OLS2 (2)	FE1 (3)	FE2 (4)
$\hat{e}^2_{\ln y}$	−.003	−.004	−.009	−.011
	(.007)	(.007)	(.010)	(.010)
LnRealEarn	.012**	.034***	.021***	.051***
	(.005)	(.011)	(.008)	(.017)
LnWorkSatis	.618***	.615***	.588***	.588***
	(.014)	(.014)	(.025)	(.025)
Married		.027**		.018
		(.010)		(.020)
Age		−.006*		−.007
		(.003)		(.013)
Age^2		.00007*		.0001
		(.00004)		(.0002)
EarnQuart = 2		−.003		−.027
		(.018)		(.026)
EarnQuart = 3		−.050**		−.068*
		(.023)		(.035)
EarnQuart = 4		−.061*		−.096
		(.032)		(.049)**
Constant	.428***	.487***	.438***	.524**
	(.019)	(.059)	(.031)	(.247)
Observations ($N \times T$)	2,507	2,507	2,507	2,507
R^2	.45	.454	.42	.423

Note. Robust standard errors are in parentheses. OLS = ordinary least squares; FE = fixed effects; EarnQuart = earnings quartile.
* Confidence: 90%.
** Confidence: 95%.
*** Confidence: 99%.

of lottery games. This complements our analysis of life satisfaction and high-lights the role of downside risk in the domain of decision-making (Kahneman, Wakker, and Sarin 1997).

The experiment, extensively described in Falco (2014), was run in 2007, with a random subsample of 307 GUHPS respondents. It consisted of 21 choices between pairs of monetary lotteries. Each "game" was framed as a choice between two opaque urns containing marbles of different colors and, correspondingly, different monetary values.[26] After being shown the composition of each urn, respondents were asked to choose the one from which they would prefer to draw a marble. Before making their choices, they were informed that at the end of the game one of their 21 preferred lotteries would be randomly selected and played out. The winnings of that game would then be paid to the respondent. Monetary incentives of this kind are used to induce truthful revelation of preferences.

---

[26] Barr (2007) has a detailed description of the experimental setup.

Choices were framed in terms of losses and gains with respect to the reference point of no gain over the initial endowment. This standard manipulation is ubiquitous in the experimental literature on loss aversion.

Using data from the experiment, we perform a maximum likelihood estimation of the following utility function, which incorporates a loss-aversion parameter, λ:[27]

$$u(x) = \begin{cases} x^\alpha & \text{if } x \geq 0, \\ -\lambda(-x^\beta) & \text{if } x < 0. \end{cases} \tag{8}$$

This is a standard parameterization of utility functions in the prospect-theory literature (Wakker 2010). An estimate of λ greater than 1 is evidence of loss aversion: losses are felt more than gains. In line with prospect theory, we further assume that prospects are evaluated as a weighted sum of the utilities of the various outcomes, where the weights are transformations of the actual probabilities given by the following probability-weighting function:

$$\omega(p) = \begin{cases} \dfrac{p^\gamma}{\left[p^\gamma + (1-p)^\gamma\right]^{1/\gamma}} & \text{if } x \geq 0, \\[2em] \dfrac{p^\phi}{\left[p^\phi + (1-p)^\phi\right]^{1/\phi}} & \text{if } x < 0. \end{cases} \tag{9}$$

Given these assumptions about the form of the utility function and the weighting of probabilities, we can calculate the difference in utility between the two lotteries in each pair. For each of the two lotteries R and L we obtain

$$\nabla PU = \frac{\sum_R \omega(p_R) u(x_R) - \sum_L \omega(p_L) u(x_L)}{\mu}, \tag{10}$$

where ∇PU includes an error term capturing the possibility that individuals make mistakes when assessing the utility they would derive from a lottery. We model this error term following the recommendations of the literature (Hey and Orme 1994; Andersen et al. 2010). The choice of a lottery is then modeled as a stochastic function of ∇PU. The log-likelihood is hence given by

$$\begin{aligned} \ln L(&\alpha, \beta, \lambda, \gamma, \phi, \mu; y, X) \\ &= \sum_i [(\ln \Phi(\nabla PU)|y_i = R) + (\ln \Phi(1 - \nabla PU)|y_i = L)]. \end{aligned} \tag{11}$$

---

[27] Following prospect theory, this utility function has separate parameters to define the curvature in the loss and gain domains. We assume that the reference point is 0 and define x as the gain over the initial endowment. This is consistent with the framing of the lotteries.

Parameter	Estimate	Robust Standard Error	z	$p > |z|$
α	.42	.017	24.24	.00
β	1.90	.074	25.77	.00
λ	1.77	.169	10.45	.00

Details on the estimation procedure are further outlined in Harrison (2008) and Falco (2014).

We first estimate the parameters of utility function (8), pooling the choices of all respondents. We cluster standard errors at the individual level, as recommended by Andersen et al. (2010). The estimate of loss aversion we obtain is 1.77 (table 5), which is in line with previous experimental findings (Booij, Praag, and Kuilen 2010; Wakker 2010). Using a standard test, we can reject the null of $\lambda = 1$ at a 1% significance level.

Furthermore, we attempt to estimate a coefficient of loss aversion for each individual in the sample. Our maximum likelihood routine converges for 266 respondents. However, for 45 of them we obtain estimates of λ above 10, which are inconsistent with the upper bounds reported in other studies (Booij, Praag, and Kuilen 2010; Wakker 2010). We exclude these from the analysis. Out of the remaining observations, we estimate λ above 1, indicating loss aversion, for 55% of individuals. The precision of these individual estimates is, however, low, so we are able to reject the null hypothesis of $\lambda = 1$ for only 22% of the respondents. Figure 7 shows the distribution of estimated loss-aversion coefficients.


Figure 7. Distribution of loss-aversion coefficients λ; Epanechikov kernel; bandwidth = 0.4553.

Overall, the results in this section complement the discussion above by showing that the workers in our sample are characterized by significant loss aversion, and therefore that downside risk will affect not only their well-being but also their decisions. Thus, labor markets that expose workers to large downside risks not only harm workers' welfare but may also lead people to forgo profitable (but risky) investment opportunities (Blumberg and Kremer 2014).

An interesting extension of the analysis would be to investigate whether individuals with higher loss aversion experience higher losses in well-being as a result of vulnerability. While it is possible, this may not necessarily be the case, since loss aversion describes the weight that people give to outcomes below a reference point when making decisions, and this "decision utility" may not correlate perfectly with "experienced utility" (Kahneman, Wakker, and Sarin 1997). However, as the individual estimates of loss aversion are rather noisy, we are not in a position to convincingly investigate in our data whether income vulnerability causes larger well-being costs for loss-averse individuals.[28]

V. Conclusions

This article investigates the relationship between income and well-being in a growing developing country, with a focus on the previously unexplored link between the risk of income poverty and happiness. Using unique longitudinal data from a representative household survey from urban Ghana, we are able to measure the probability of income poverty at the individual level and explore its relationship with life satisfaction. Our results are compelling.

First, our analysis reveals a substantial risk of poverty for a large share of the population. Second, we find a strong negative relationship between vulnerability to poverty and life satisfaction, over and above the positive income effect commonly documented in the literature. The result is both statistically significant and economically meaningful. Reducing vulnerability by 20 percentage points has the same effect on well-being as increasing income by 50%. Interestingly, we find that failing to control for individual fixed effects leads to significant bias and misleading conclusions. Further, we attempt to disentangle

---

[28] There is a second reason not to attempt this analysis. In the experiment we measure loss aversion with respect to a reference point created by a fixed endowment, which represents the "status quo" before the decision. This is standard practice in the literature that studies loss aversion. In the previous sections, on the other hand, we have studied the risk of falling below a poverty line. The poverty line will differ from the status quo for many individuals. Thus, while we study attitudes toward downside risk in both cases, the two measures we use define "downside risk" differently. In general, the literature on loss aversion has not yet fully understood how individuals choose reference points, and we flag this as an area for future research (Kőszegi and Rabin 2006). We thank an anonymous reviewer for raising this point.

the effect of downside risk on happiness from the effect of two-sided uncertainty. We find that the former has the clearer impact on subjective well-being. Finally, in a matched behavioral experiment that elicits respondents' attitudes toward risky prospects, we find evidence of significant levels of loss aversion among our respondents. This suggests that the effect of downside risk is not limited to life evaluation but extends to decision making.

Our results highlight the importance of social protection as an instrument to shield workers from the risk of falling into poverty. Loss-averse agents will be more willing to undertake productive investments when safety nets and insurance minimize this risk. Importantly, our evidence suggests that these policies will directly improve life satisfaction and should be considered alongside interventions that target other determinants of well-being, such as working conditions and work-life balance. More broadly, our results suggest that the well-being gains that developing countries can obtain by raising average incomes may be limited if the risk of falling into poverty remains high for a large share of the population.

References

Andersen, Steffen, Glenn W. Harrison, Morten Igel Lau, and E. Elisabet Rutström. 2010. "Behavioral Econometrics for Psychologists." *Journal of Economic Psychology* 31, no. 4:553–76.

Banerjee, Abhijit V., and Esther Duflo. 2007. "The Economic Lives of the Poor." *Journal of Economic Perspectives* 21, no. 1:141–68.

Barr, Abigail. 2007. "Attitudes to Risk in Ghana: Field Manual." Unpublished manuscript, University of Oxford.

Benz, Matthias, and Bruno S. Frey. 2008. "The Value of Doing What You Like: Evidence from the Self-Employed in 23 Countries." *Journal of Economic Behavior and Organization* 68, no. 3–4:445–55.

Blanchflower, David G., and Andrew J. Oswald. 2002. "Well-Being over Time in Britain and the USA." *Journal of Public Economics* 88, no. 7–8:1359–86.

Blumberg, Seth, and Michael Kremer. 2014. "Dynamic Loss Aversion, Growth, and Development." Unpublished manuscript.

Bocquier, Philippe, Christophe J. Nordman, and Aude Vescovo. 2010. "Employment Vulnerability and Earnings in Urban West Africa." *World Development* 38, no. 9: 1297–1314.

Booij, Adam, Bernard Praag, and Gijs Kuilen. 2010. "A Parametric Analysis of Prospect Theory's Functionals for the General Population." *Theory and Decision* 68, no. 1: 115–48.

Chaudhuri, Shubham. 2003. "Assessing Vulnerability to Poverty: Concepts, Empirical Methods and Illustrative Examples." Unpublished manuscript, Department of Economics, Columbia University.

Chaudhuri, Shubham, Jyotsna Jalan, and Asep Suryahadi. 2002. "Assessing Household Vulnerability to Poverty from Cross-Sectional Data: A Methodology and Es-

timates from Indonesia." Discussion Paper no. 0102-52, Department of Economics, Columbia University.

Clark, Andrew E., Paul Frijters, and Michael A. Shields. 2008. "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles." *Journal of Economic Literature* 46, no. 1:95–144.

De Neve, Jan-Emmanuel, and Andrew J. Oswald. 2012. "Estimating the Influence of Life Satisfaction and Positive Affect on Later Income Using Sibling Fixed-Effects." CEP Discussion Paper no. 1176 (November), Centre for Economic Performance, London School of Economics.

De Neve, Kristina M., and Harris Cooper. 1998. "The Happy Personality: A Meta-analysis of 127 Personality Traits of Subjective Well Being." *Psychological Bulletin* 125:197–229.

Dercon, Stefan. 1996. "Risk, Crop Choice, and Savings: Evidence from Tanzania." *Economic Development and Cultural Change* 44, no. 3:485–514.

———. 2004. "Growth and Shocks: Evidence from Rural Ethiopia." *Journal of Development Economics* 74, no. 2:309–29.

Dercon, Stefan, and Luc Christiaensen. 2007. "Consumption Risk, Technology Adoption and Poverty Traps: Evidence from Ethiopia." CSAE Working Paper no. 2007-06, Centre for the Study of African Economies, Oxford.

Dercon, Stefan, John Hoddinott, and Tassew Woldehanna. 2005. "Shocks and Consumption in 15 Ethiopian Villages, 1999–2004." *Journal of African Economies* 14, no. 4:559–85.

Di Tella, Rafael, John Haisken-De New, and Robert MacCulloch. 2007. "Happiness Adaptation to Income and to Status in an Individual Panel." NBER Working Paper no. 13159 (June), National Bureau of Economic Research, Cambridge, MA.

Easterlin, Richard A. 2001. "Income and Happiness: Towards a Unified Theory." *Economic Journal* 111, no. 473:465–84.

Fafchamps, Marcel. 2003. *Rural Poverty, Risk, and Development*. Cheltenham: Elgar.

———. 2009. "Vulnerability, Risk Management, and Agricultural Development." Paper presented at the AERC (African Economic Research Consortium) Conference on Agriculture and Development, Mombasa, Kenya, May 28–30.

Fafchamps, Marcel, David McKenzie, Simon Quinn, and Christopher Woodruff. 2012. "Using PDA Consistency Checks to Increase the Precision of Profits and Sales Measurement in Panels." *Journal of Development Economics* 98, no. 1:51–57.

Fafchamps, Marcel, and Forhad Shilpi. 2008. "Subjective Welfare, Isolation, and Relative Consumption." *Journal of Development Economics* 86, no. 1:43–60. doi:10.1016/j.jdeveco.2007.08.004.

———. 2009. "Isolation and Subjective Welfare: Evidence from South Asia." *Economic Development and Cultural Change* 57, no. 4:641–83.

Falco, Paolo. 2014. "Does Risk Matter for Occupational Choices? Experimental Evidence from an African Labour Market." *Labour Economics* 28:96–109.

Falco, Paolo, and Luke Haywood. 2016. "Entrepreneurship versus Joblessness: Explaining the Rise in Self-Employment." *Journal of Development Economics* 118: 245–65.

Falco, Paolo, Andrew Kerr, Neil Rankin, Justin Sandefur, and Francis Teal. 2011. "The Returns to Formality and Informality in Urban Africa." *Labour Economics* 18, suppl. 1:S23–S31.

Falco, Paolo, William F. Maloney, Bob Rijkers, and Mauricio Sarrias. 2015. "Heterogeneity in Subjective Wellbeing: An Application to Occupational Allocation in Africa." *Journal of Economic Behavior and Organization* 111:137–53.

Ferrer-i-Carbonell, Ada, and Paul Frijters. 2004. "How Important Is Methodology for the Estimates of the Determinants of Happiness?" *Economic Journal* 114, no. 497: 641–59.

Frey, Bruno S., and Alois Stutzer. 2002. "What Can Economists Learn from Happiness Research?" *Journal of Economic Literature* 40, no. 2:402–35.

Graham, Carol, Andrew Eggers, and Sandip Sukhtankar. 2004. "Does Happiness Pay? An Exploration Based on Panel Data from Russia." *Journal of Economic Behavior and Organization* 55, no. 3:319–42.

Harrison, Glenn W. 2008. "Maximum Likelihood Estimation of Utility Functions Using *Stata*." Working Paper no. 06-12, Department of Economics, University of Central Florida.

Hey, John D., and Chris Orme. 1994. "Investigating Generalizations of Expected Utility Theory Using Experimental Data." *Econometrica* 62, no. 6:1291–326.

Howell, Ryan T., and Colleen J. Howell. 2008. "The Relation of Economic Status to Subjective Well-Being in Developing Countries: A Meta-analysis." *Psychological Bulletin* 134, no. 4:536–60.

Kahneman, Daniel, and Angus Deaton. 2010. "High Income Improves Evaluation of Life but Not Emotional Well-Being." *Proceedings of the National Academy of Sciences of the USA* 107, no. 38:16489–93.

Kahneman, Daniel, Peter P. Wakker, and Rakesh Sarin. 1997. "Back to Bentham? Explorations of Experienced Utility." *Quarterly Journal of Economics* 112, no. 2: 375–405.

Kingdon, Geeta, and John Knight. 2004. "Community, Comparisons and Subjective Well-Being in a Divided Society." CSAE Working Paper no. 2004-21, Centre for the Study of African Economies, Oxford.

Knight, John, and Ramani Gunatilaka. 2008. "Aspirations, Adaptation and Subjective Well-Being of Rural-Urban Migrants in China." Economics Working Paper no. 381 (January), University of Oxford.

Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121, no. 4:1133–65.

Layard, Richard. 2005. *Happiness: Lessons from a New Science*. London: Penguin.

Ligon, Ethan, and Laura Schechter. 2003. "Measuring Vulnerability." *Economic Journal* 113, no. 486:C95–C102.

———. 2004. "Evaluating Different Approaches to Estimating Vulnerability." Social Protection Discussion Paper no. 0410 (June), World Bank, Washington, DC.

Luttmer, Erzo F. P. 2005. "Neighbours as Negatives: Relative Earnings and Well-Being." *Quarterly Journal of Economics* 120, no. 3:963–1002.

Meghir, Costas, and Luigi Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity." *Econometrica* 72, no. 1:1–32.

Mullainathan, Sendhil, and Marianne Bertrand. 2001. "Do People Mean What They Say? Implications for Subjective Survey Data." *American Economic Review* 91, no. 2: 67–72.

Nsowah-Nuamah, Nicholas, Francis Teal, and Moses Awoonor-Williams. 2010. "Jobs, Skills and Incomes in Ghana: How Was Poverty Halved?" CSAE Working Paper no. 2010-01, Centre for the Study of African Economies, Oxford.

Oswald, Andrew J., and Stephen Wu. 2010. "Objective Confirmation of Subjective Measures of Human Well-Being: Evidence from the USA." IZA Discussion Paper no. 4695 (January), Institute for the Study of Labor, Bonn.

Powdthavee, Nattavudh. 2010. "How Much Does Money Really Matter? Estimating the Causal Effects of Income on Happiness." *Empirical Economics* 39, no. 1:77–92.

Rankin, Neil, Justin Sandefur, and Francis Teal. 2010. "Learning and Earning in Africa: Where Are the Returns to Education High?" CSAE Working Paper no. 2010-02, 2010, Centre for the Study of African Economies, Oxford.

Rosenzweig, Mark R., and Hans P. Binswanger. 1993. "Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments." *Economic Journal* 103, no. 416:56–78.

Wakker, Peter P. 2010. *Prospect Theory*. Cambridge: Cambridge University Press.